Virtual worlds are laboratories for agency

Virtual Worlds in the Age of AI



State • dynamics • agency • consequences

From representation to situated action

Which of these six systems is a world?

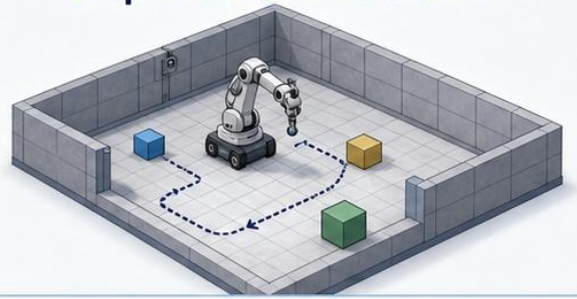
1. static 3D museum



2. text adventure



3. episodic robot simulator



4. persistent social sandbox



5. sensor-synchronised factory



6. phone AR maintenance experience



- ○ WORLD
- ○ ENVIRONMENT
- ○ INSUFFICIENT EVIDENCE

A beautiful scene is not yet a world

SCENE

A scene describes what exists



CAUSAL TRANSITION

action

transition

consequence

WORLD

A world defines what can happen



EVENT TRACE

00:00 **ACTION**
Patient takes a step

00:01 **TRANSITION**
Door opens (state change)

00:02 **CONSEQUENCE**
More light enters room (state change)

00:03 **CONSEQUENCE**
Patient is now closer to target (state change)

A scene describes what exists

A world defines what can happen

Photorealism does not create consequences

Labels are engineering commitments



- Architecture
- Evidence
- Evaluation
- Governance

The classical social definition



Bell (2008)

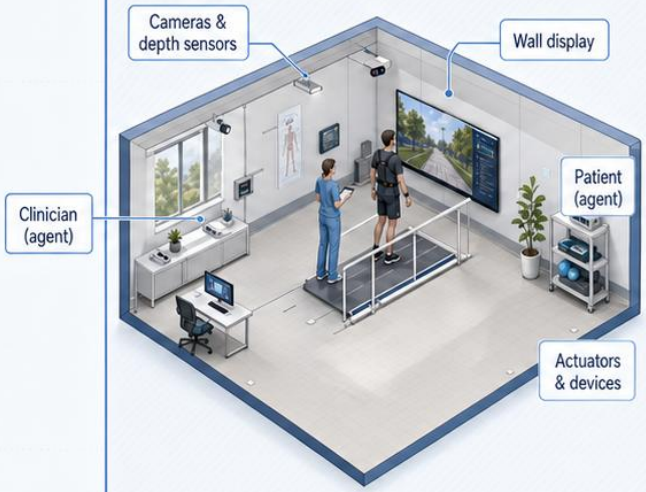
AI environments stress the boundary

The classical social-world boundary still matters—and the course broadens it for AI.

- **Episodic**
- **Single-agent**
- **Robotic**
- **Synthetic**
- **Generative**

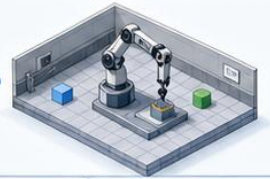
CLASSICAL SOCIAL-WORLD BOUNDARY

People, avatars, synchrony, persistence, and networks are central.



Persistent, relational, social interaction.

AI ENVIRONMENTS BEYOND THE CLASSICAL BOUNDARY



EPISODIC ROBOT SIMULATOR

Short episodes with reset.
State, dynamics, consequences.



SINGLE SOFTWARE-AGENT WORLD

One agent, discrete state space.
Perception, action, reward.



SYNTHETIC-DATA FACTORY

Generate images, LiDAR, trajectories, labels.
No human participants.



GENERATIVE WORLD MODEL

Learned dynamics rollouts.
Counterfactuals and planning.
Useful but not inhabited.



SYNTHETIC EVALUATION WORLD

Benchmarks and stress tests.
Deterministic or stochastic.
Designed for measurement.

BROADER COURSE SCOPE

The course working definition



A world is an evolving state

A transition system for virtual worlds

- S_t: world state
- A_t: agent actions
- E_t: external events
- $S_{t+1} = F(S_t, A_t, E_t)$

$$S_{t+1} = F(S_t, A_t, E_t)$$



Key components

State S_t
Everything that makes a difference

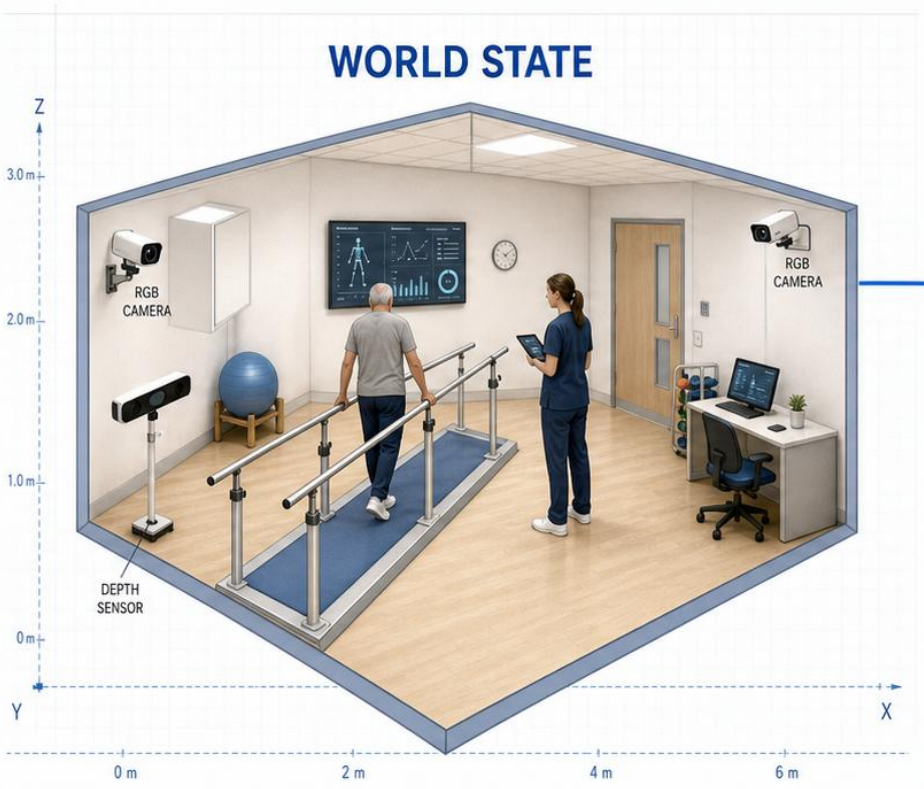
Actions A_t
Agent's choices and controls

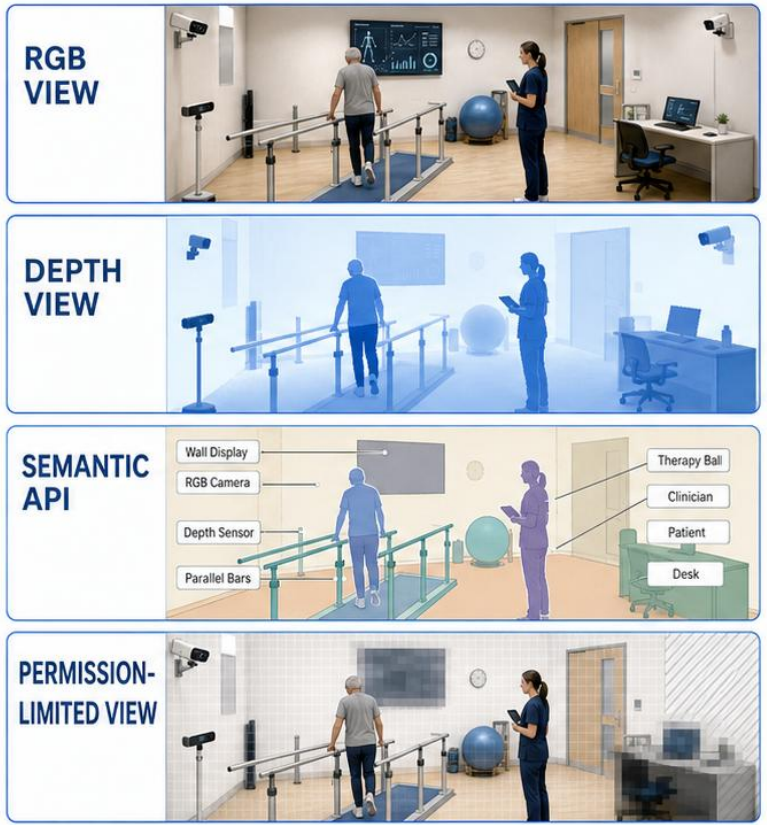
Events E_t
External happenings, sensor updates

Time
Discrete or continuous

The world evolves through consequences.

The world knows more than an agent observes





Observation depends on viewpoint, sensors, embodiment, and permission.

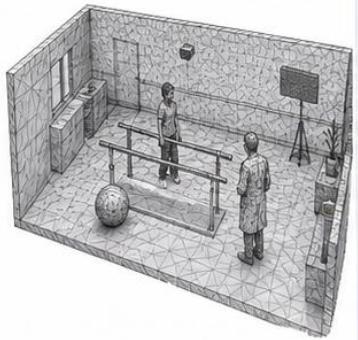
Representation shapes storage and sensing



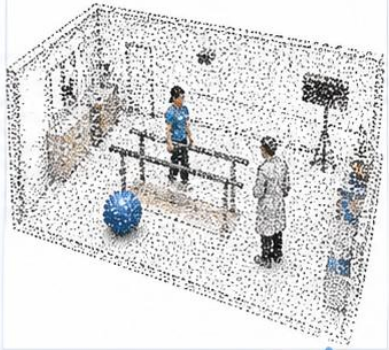
CONSTANT VIEWPOINT

Same camera and rehabilitation-room geometry across all representations.

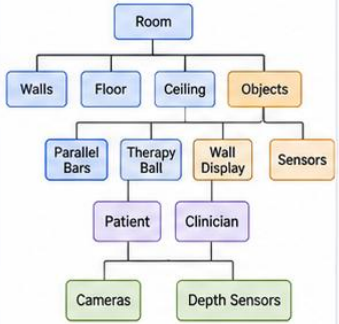
1 MESH



2 POINT CLOUD



3 SCENE GRAPH



4 VOXEL GRID



5 SEMANTIC MAP



STORE
How the data is saved and organized.

RENDER
How the room is drawn or visualized.

QUERY
How we find, reason, and detect.

SENSE
How sensors capture and interpret.

Time is a process, not decoration

WALL TIME
real clock



00:10

SIMULATION TIME
model clock



00:25 2.5×

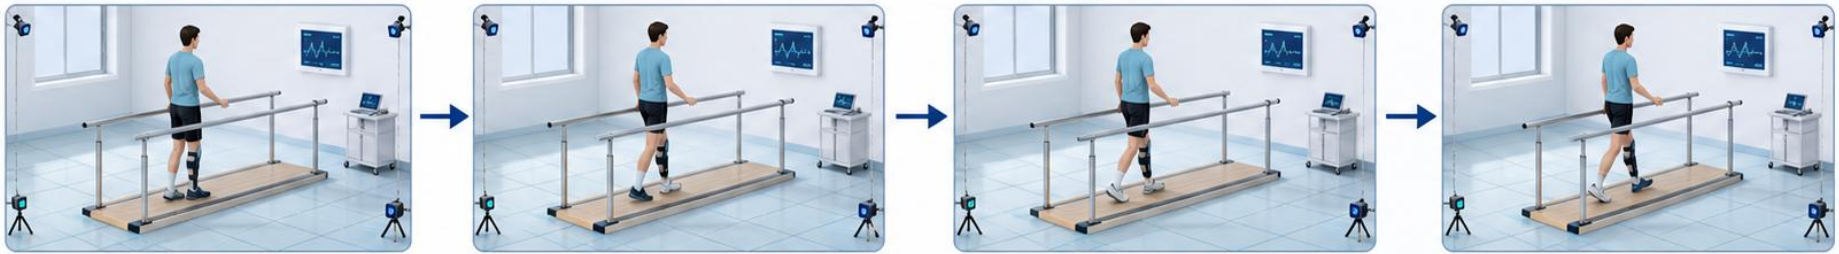
EPISODE TIME
local episode clock



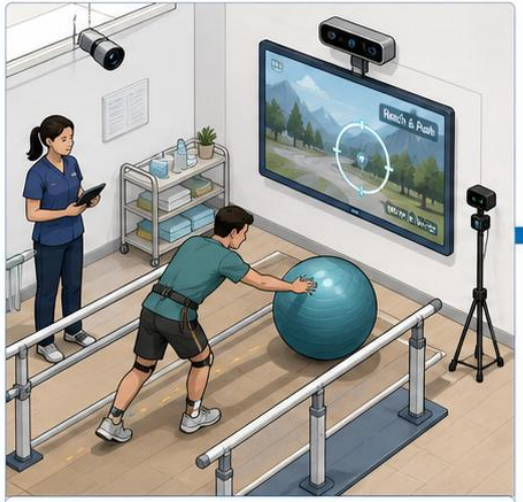
00:05

reset horizon

RESET • RATE • HORIZON



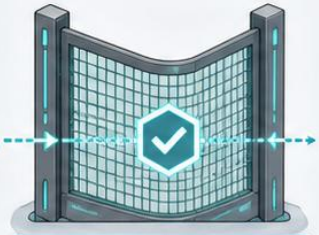
Rules turn actions into consequences



1 ATTEMPT

patient or embodied agent reaches to move the therapy ball.

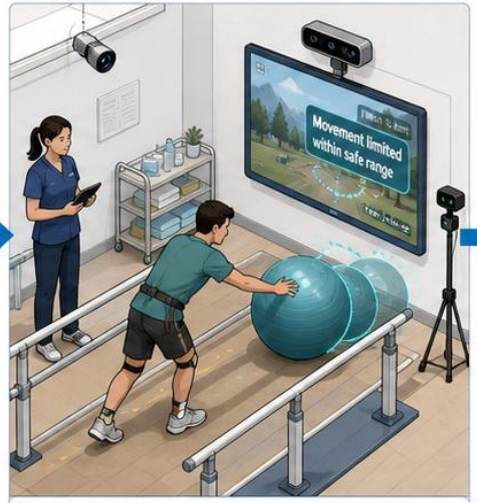
RULE ENGINE
constraint evaluation



Joint angle limit	OK
Workspace boundary	LIMIT
Therapist override	OK
Fatigue model	ADAPT

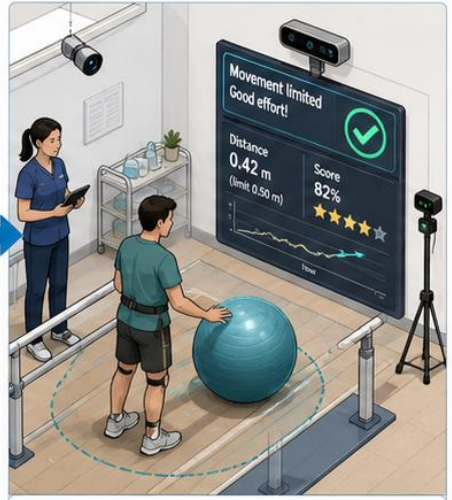
2 CONSTRAINT CHECK

the attempted action passes through a clear technical rule gate that evaluates constraints.



3 STATE TRANSITION

the world state changes according to the rule/dynamics.



4 OBSERVABLE CONSEQUENCE

show a visibly different resulting state that the agent can observe, such as ball motion being limited or modified and a clear response on the wall display.

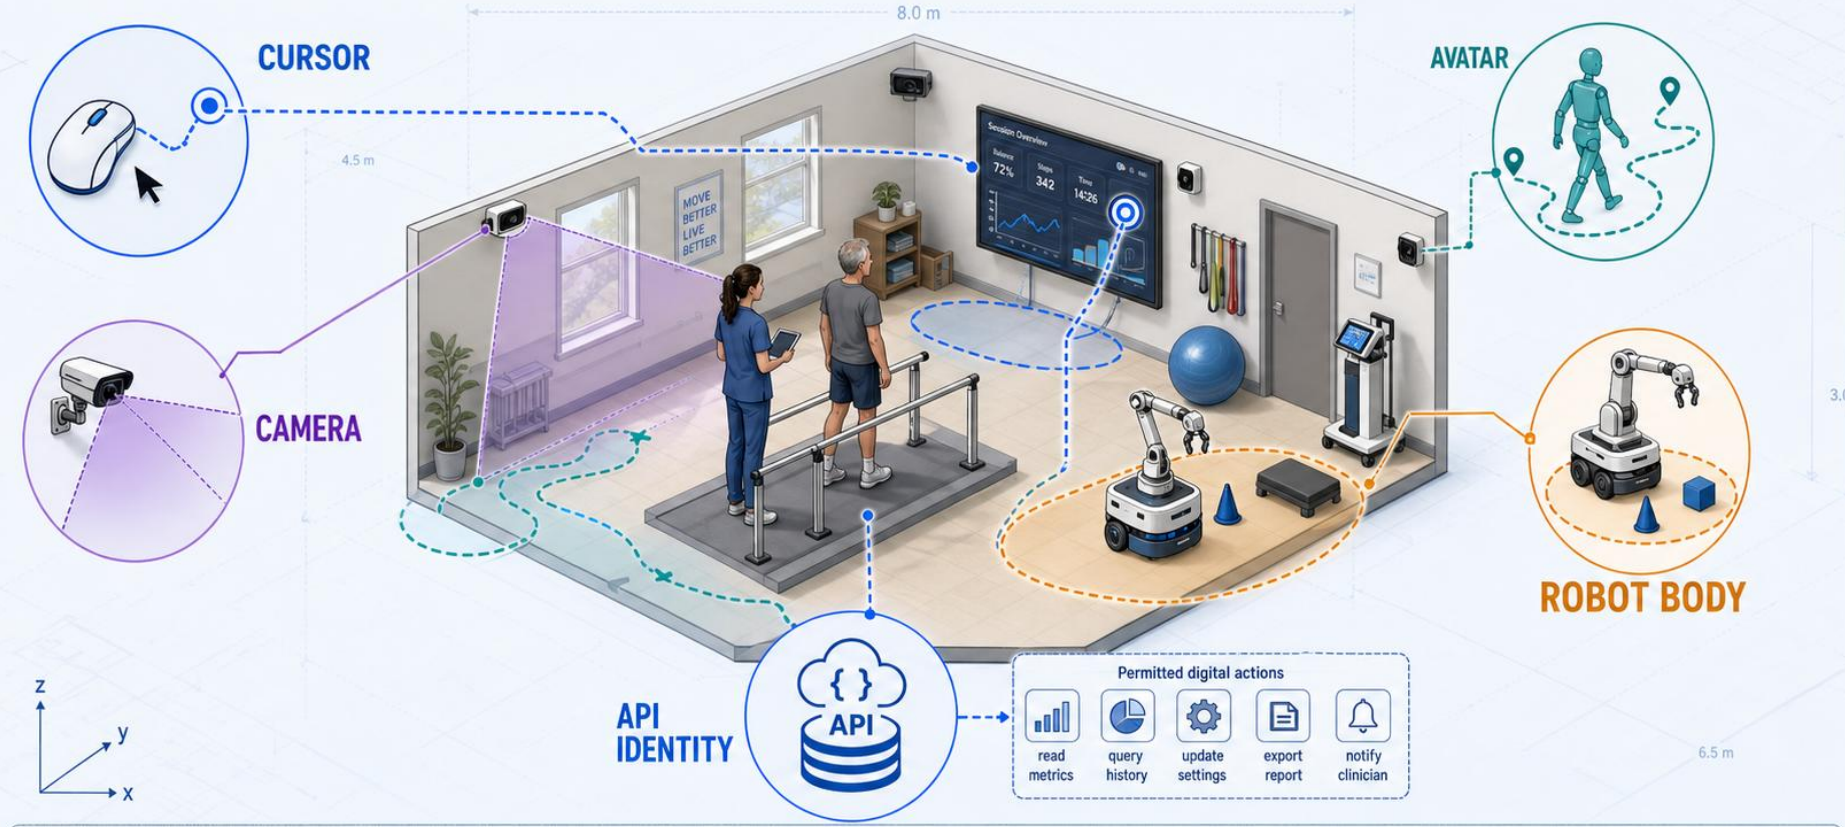
PHYSICAL

LOGICAL

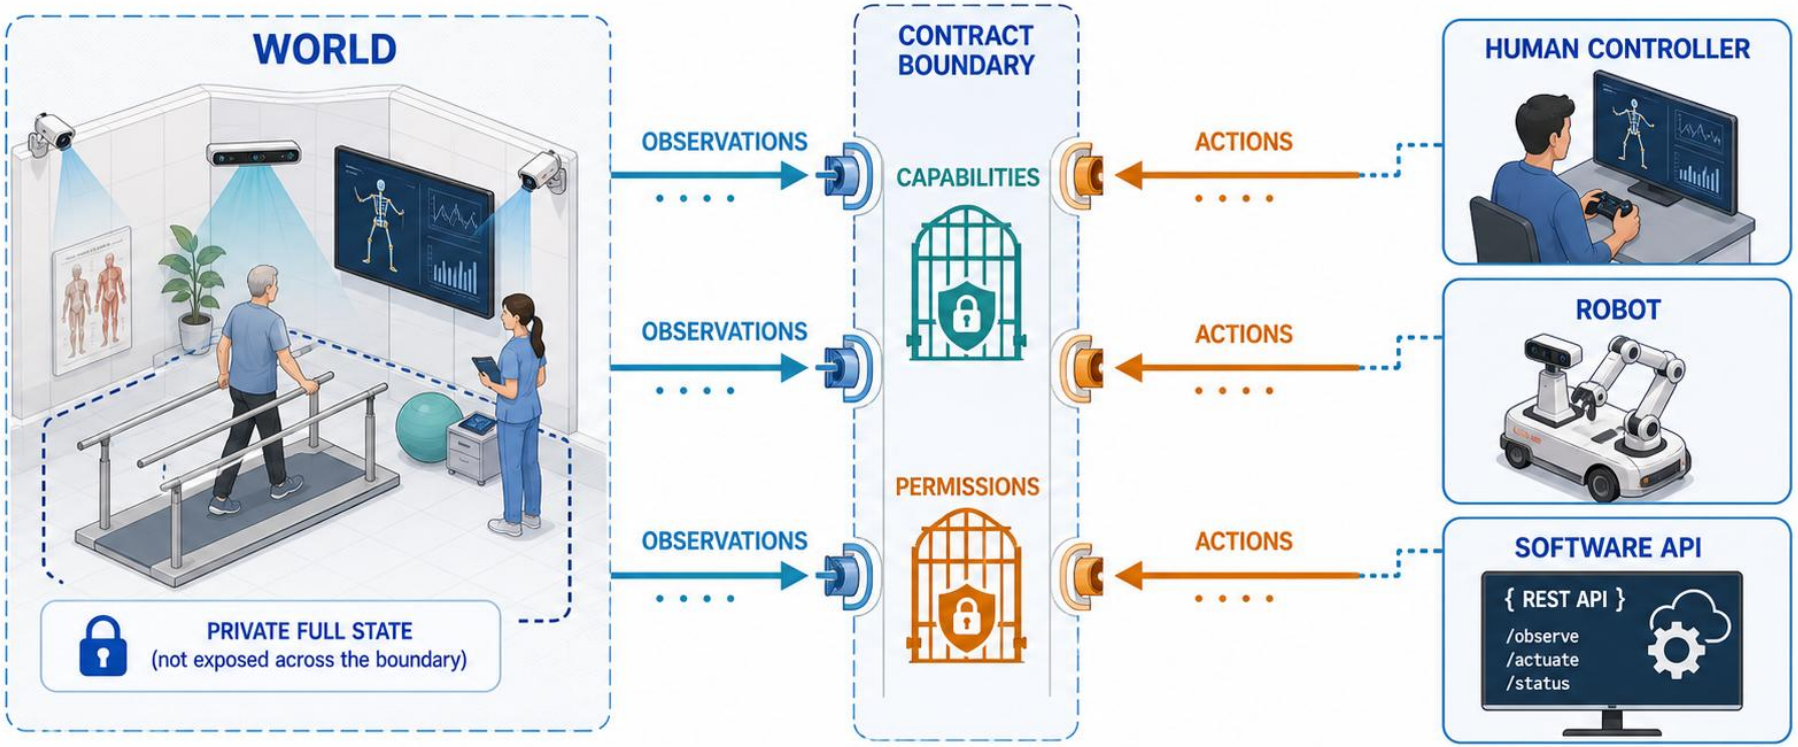
SOCIAL

LEARNED

Embodiment defines a point of intervention

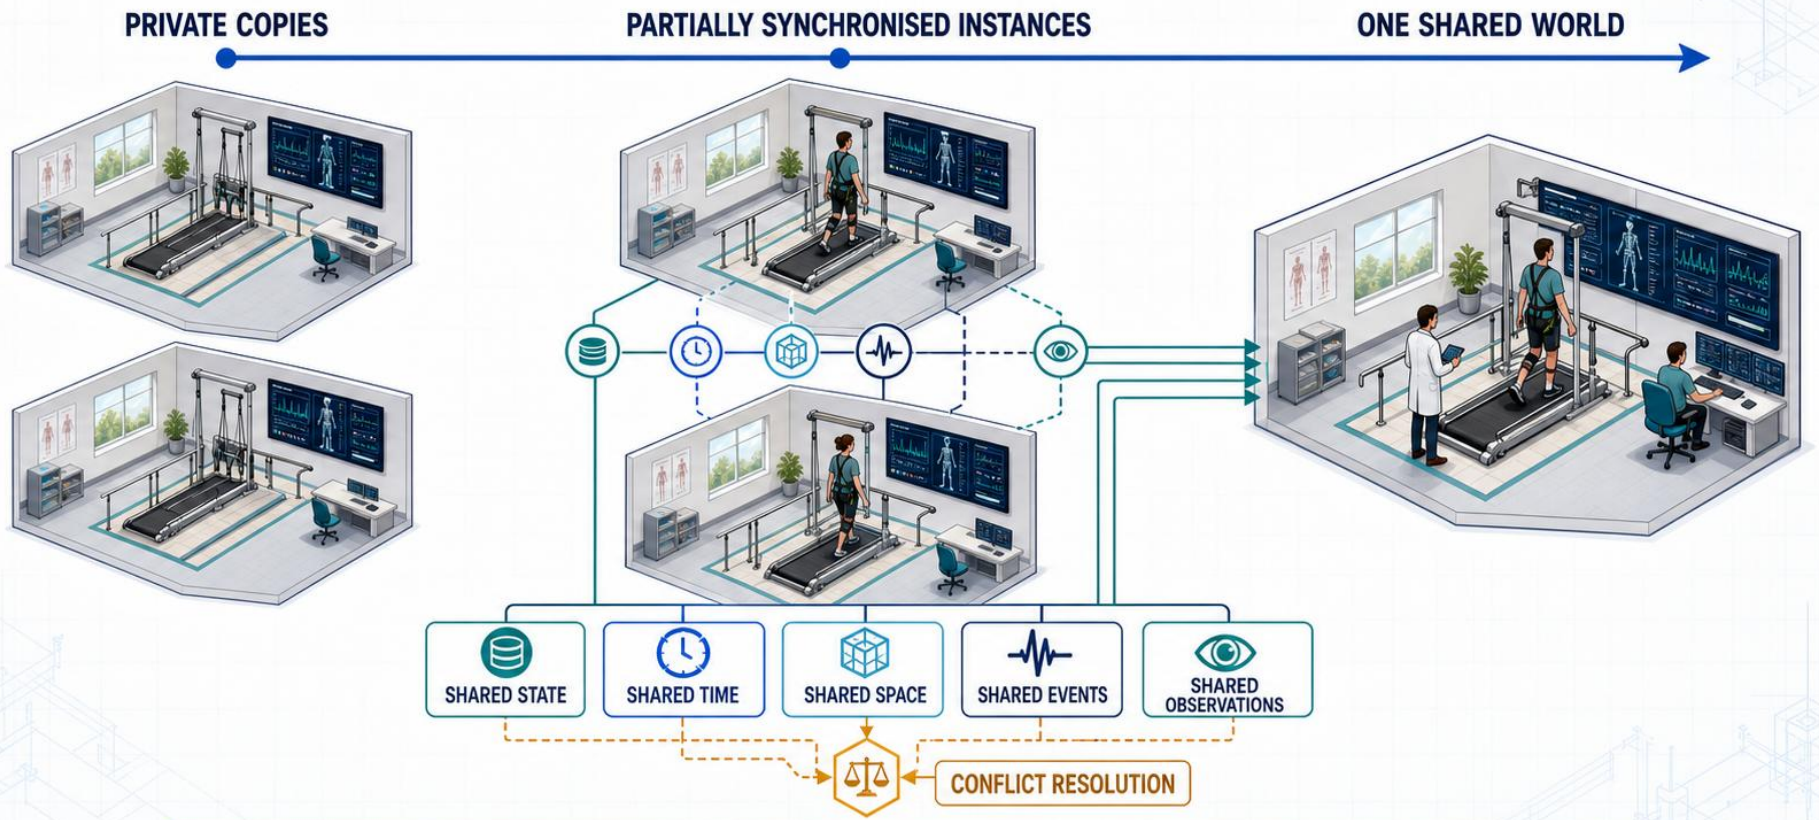


Interfaces are observation-action contracts



Interfaces expose, hide, discretise, or distort world structure.

Sharedness has several forms



CO-PRESENCE · PERSISTENCE · AUTHORITY · CONSISTENCY are separate design choices.

Physical coupling is a design variable

PHYSICAL ROOM



1 INDEPENDENT

2 MEASUREMENTS IN

10 Hz
confidence 95%

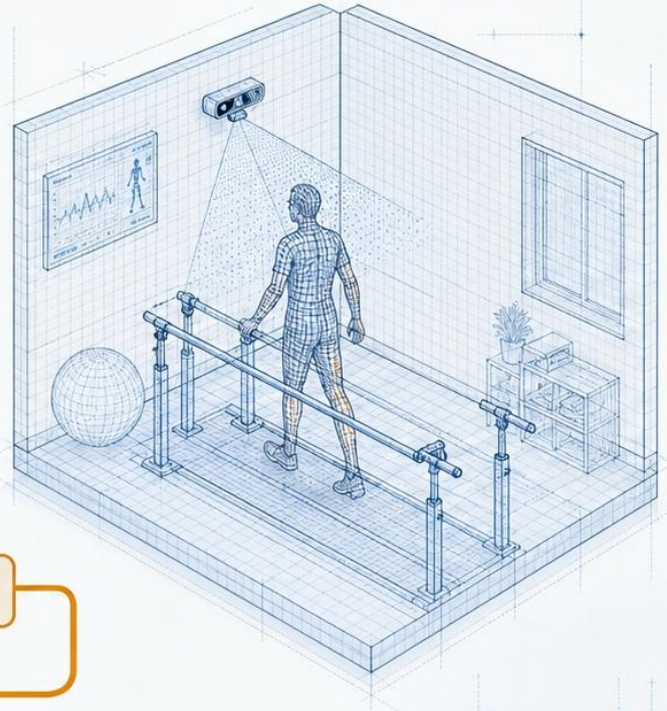
3 COMMANDS OUT

4 BIDIRECTIONAL

PERCEPTION → MODEL → ACTION

FEEDBACK

DIGITAL MODEL



Physical linkage alone does not establish digital-twin validity.

Eight dimensions replace a brittle checklist

SUBSTRATE

- representation
- state and time
- dynamics



AGENT–WORLD RELATION

- agency and embodiment
- observation and affordance
- sharedness

SYSTEM–REALITY RELATION

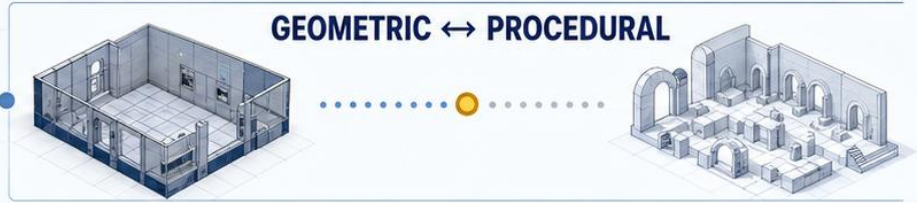
- physical coupling
- purpose and governance

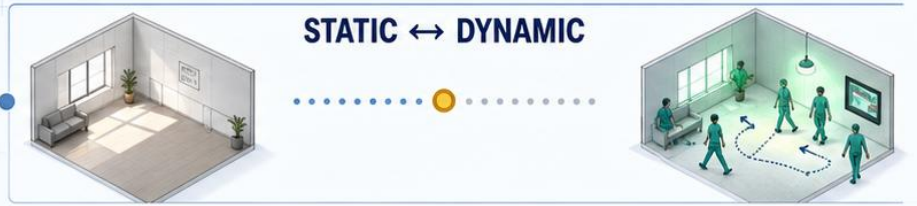
Compare — do not score

Four implementation axes

How virtual worlds are built











Useful systems are usually hybrid.

Content can be stored explicitly—or generated from rules

CONTENT CONSTRUCTION

GEOMETRIC ↔ PROCEDURAL

- **Explicit:** meshes · points · surfaces · volumes · attributes
- **Generated:** rules · parameters · grammars · models

EXPLICIT (ACQUIRED OR AUTHORED)

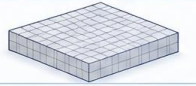
MESHES



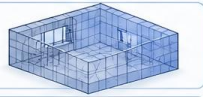
POINTS



SURFACES



VOLUMES



ATTRIBUTES



HYBRID PRODUCTION WORLD

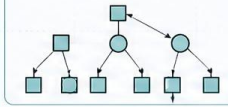
Acquired structure + procedural content





PROCEDURAL (RULES AND PARAMETERS)

RULES / GRAMMARS



```
IF window_near_wall
THEN place_plant
ELSE place_art
END IF
```

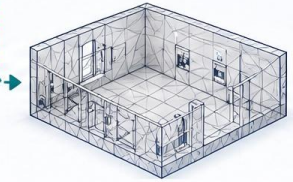
PARAMETERS

density	0.72
variety	0.45
obstacle_rate	0.30
plant_height	1.20

GENERATED OUTPUT (NOT INHERENTLY EXPLICIT)



MATERIALISE AS EXPLICIT GEOMETRY

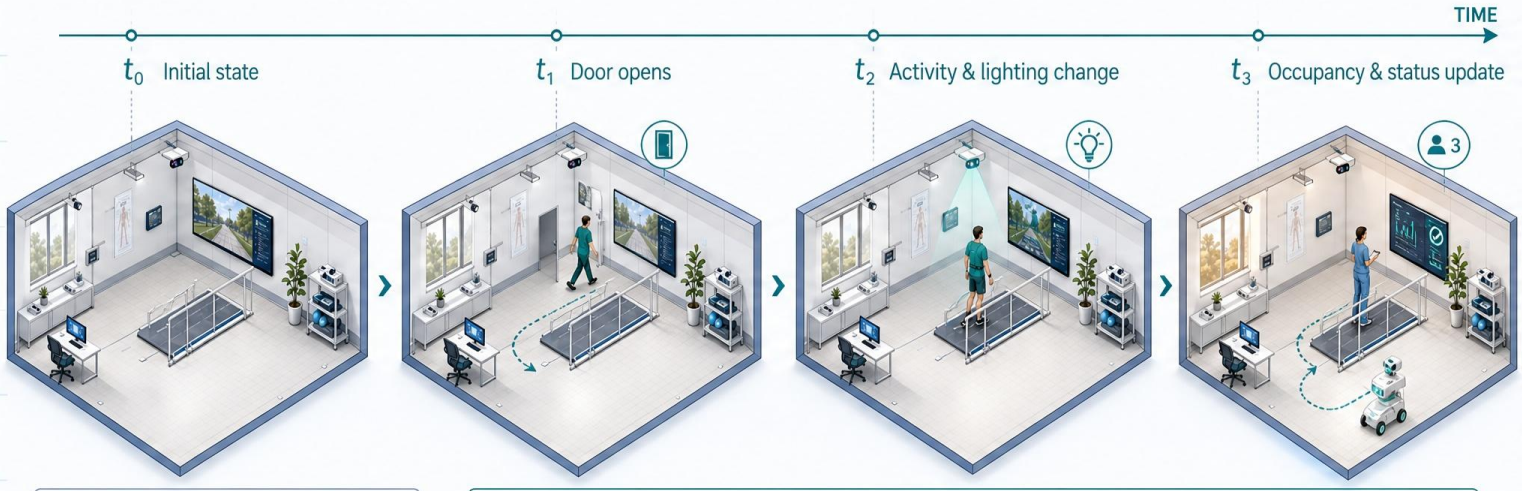


Most production worlds combine authored or acquired geometry with procedural variation.

Static geometry can host a dynamic world

▶ TEMPORAL BEHAVIOUR

STATIC ↔ DYNAMIC



FIXED SHELL (UNCHANGING)

Walls · floor · ceiling · windows
doors · fixtures · structure

MUTABLE STATE (CHANGES OVER TIME)

Door status · People & robots · Illumination · Occupancy · Object states · Semantic state & data

FROZEN REFERENCE



Fixed shell snapshot
(unchanging geometry)

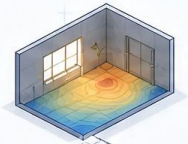
Classify the state that matters—not only the background mesh.

Real-time is a deadline—not a synonym for quality

COMPUTATION TIME

PRE-COMPUTED ↔ REAL-TIME

Offline preparation (before interaction)



Baked lightmap



Visibility map



Navigation graph



Cached trajectories

FRAME / CONTROL DEADLINE



Live computation (during interaction)

Sensing

Decisions

Physics

Shadows

Interaction

Cache can become stale. Validate and refresh.

Prepare what can remain invariant; recompute what must respond.

- **Offline:** lighting · visibility · navigation · cached trajectories
- **Live:** sensing · decisions · physics · shadows · interaction

Discrete plans can control continuous motion

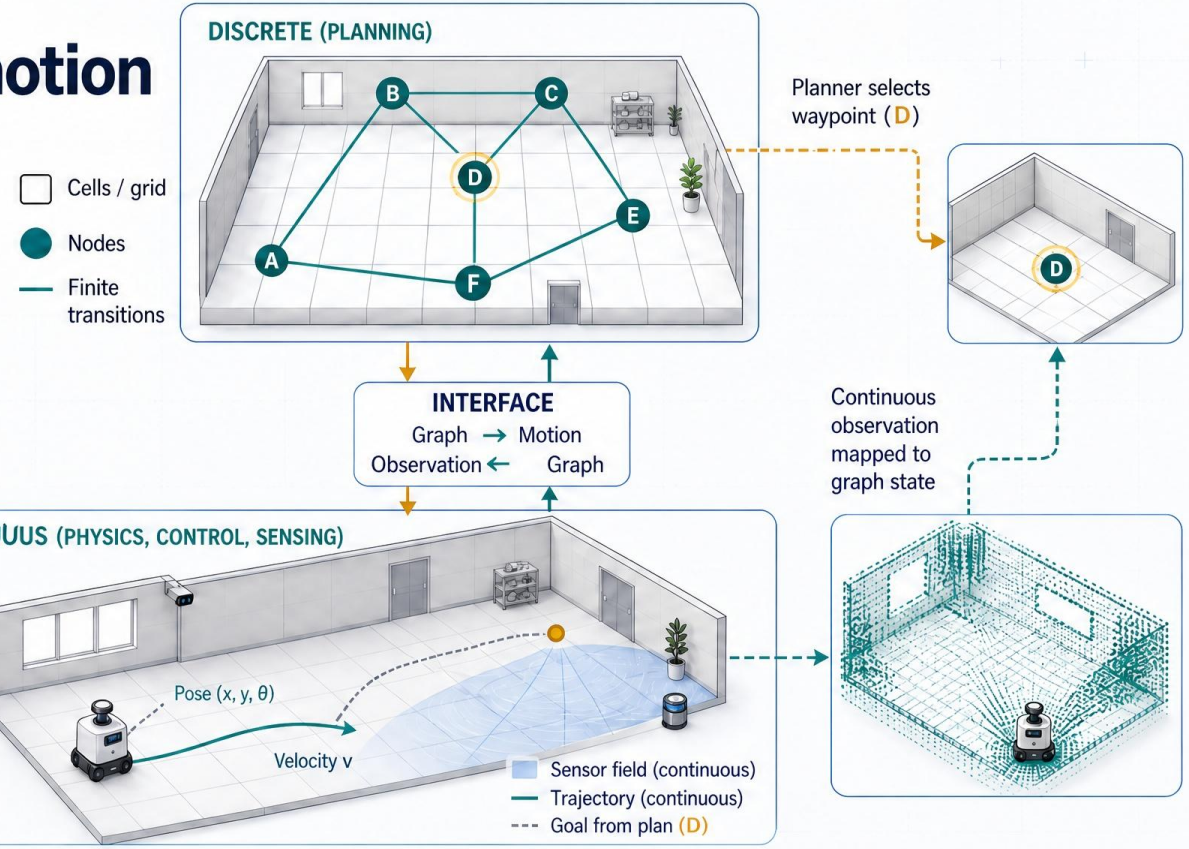
MATHEMATICAL SPACE

DISCRETE ↔ CONTINUOUS

Discrete: cells · nodes · symbols · finite actions

Continuous: coordinates · fields · poses · velocities

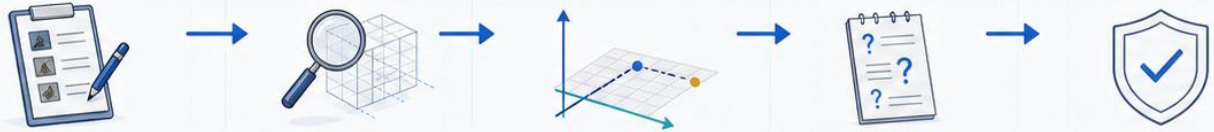
Representation may change by layer: planning, physics, sensing, and control.



This choice previews tabular versus continuous-control methods in reinforcement learning.

Classify one system and defend the boundary

1. Choose one opening scenario
2. Inspect four dimensions
3. Locate two implementation axes
4. State missing evidence
5. Defend one measurable criterion



EVIDENCE BOARD

TERM	DIMENSIONS	AXES	MISSING EVIDENCE	CRITERION

CHOOSE ONE OPENING SCENARIO

1. static 3D museum



2. text adventure



3. episodic robot simulator



4. persistent social sandbox



5. sensor-synchronised factory



6. phone AR maintenance experience



08:00

Disagreement reveals dimensions

GEOMETRIC ↔ PROCEDURAL

PRE-COMPUTED ↔ REAL-TIME

CLASSIFICATION PATH A

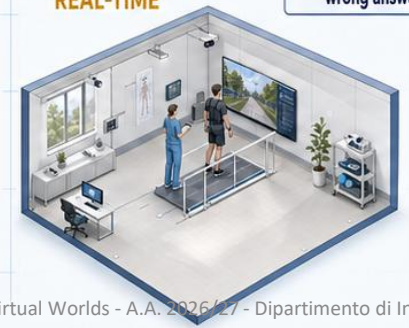
REVISION

CLASSIFICATION PATH B

STATIC ↔ DYNAMIC

DISCRETE ↔ CONTINUOUS

EVIDENCE	GEOMETRIC ↔ PROCEDURAL	STATIC ↔ DYNAMIC	PRE-COMPUTED ↔ REAL-TIME	DISCRETE ↔ CONTINUOUS	GEOMETRIC ↔ PROCEDURAL	STATIC ↔ DYNAMIC	PRE-COMPUTED ↔ REAL-TIME	DISCRETE ↔ CONTINUOUS
No transition	○	STATIC	PRE-COMPUTED	DISCRETE	GEOMETRIC	○	PRE-COMPUTED	DISCRETE
Episodic reset	PROCEDURAL	DYNAMIC	REAL-TIME	DISCRETE	PROCEDURAL	DYNAMIC	REAL-TIME	DISCRETE
Absent agency	PROCEDURAL	○	PRE-COMPUTED	DISCRETE	GEOMETRIC	STATIC	PRE-COMPUTED	DISCRETE
Partial physical world	GEOMETRIC	○	○	CONTINUOUS	GEOMETRIC	○	REAL-TIME	CONTINUOUS
Missing evidence ≠ wrong answer	○	○	○	○	○	○	○	○



EVIDENCE ≠ CRITERION

Disagreement reveals which dimension matters.



Several lineages converged

FOUNDATIONS → NETWORKED ERA → CONTEMPORARY

Interactive simulation



MUDs and social worlds



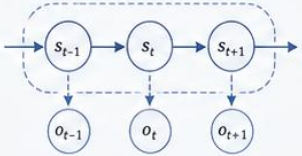
Networked 3D



XR and robotics



AI world models



INTERACTION

PERSISTENCE

NETWORKS

EMBODIMENT

PREDICTION

CONTEMPORARY VIRTUAL-WORLD SYSTEMS

CHANGING PROPERTIES

AGENCY · COUPLING · PERSISTENCE · SCALE · SOCIALITY

A game is not synonymous with a world



GAME LAYER

GAME: goals and incentives

WORLD LAYER

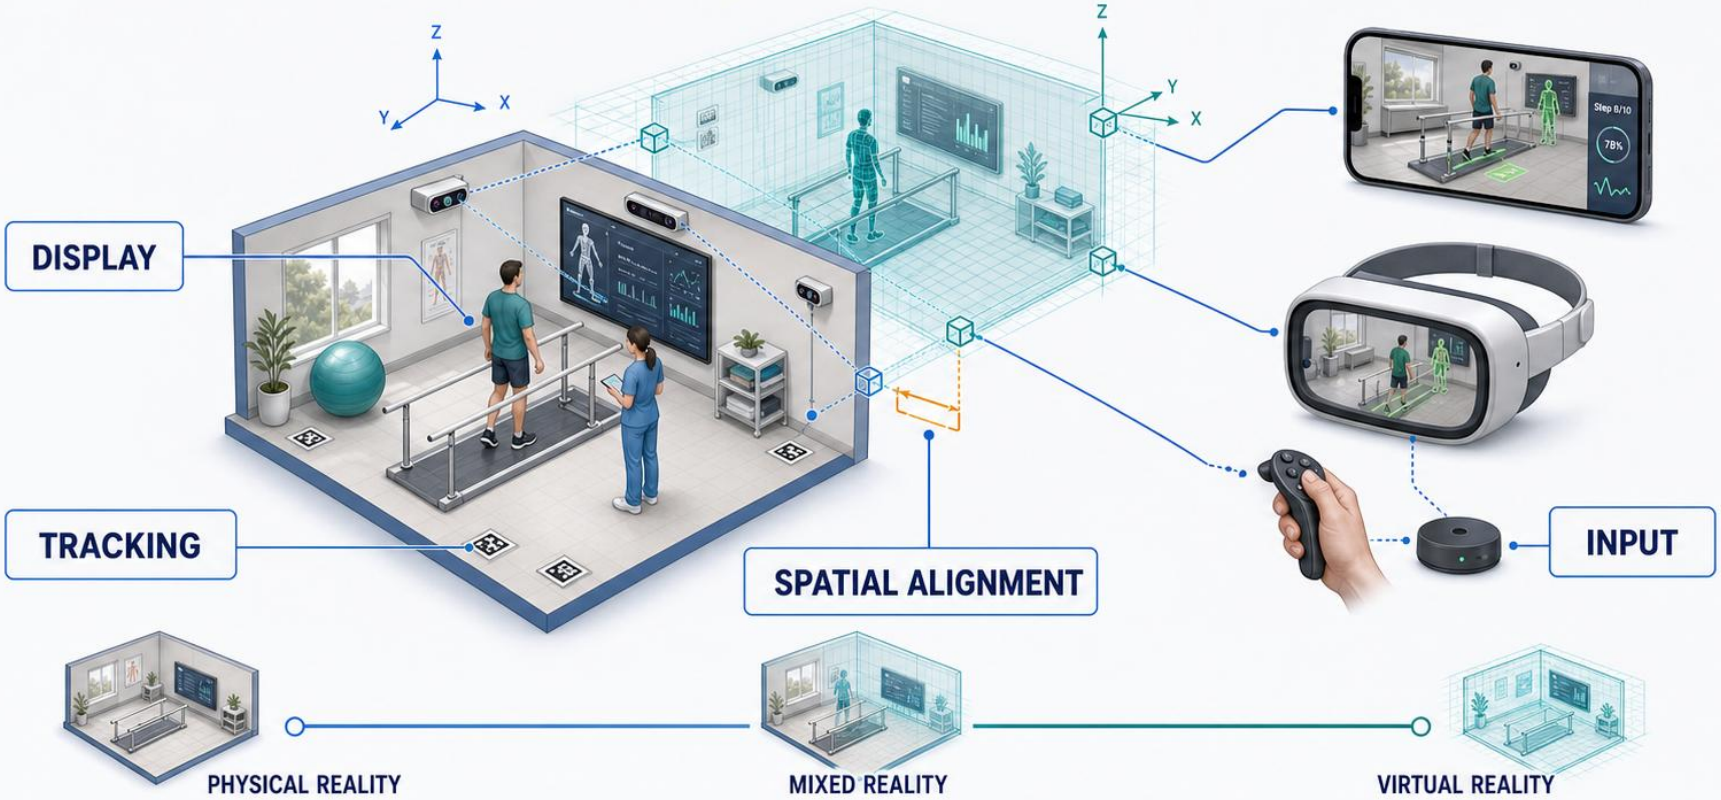
WORLD: state, space, consequences

- State
- Space
- Consequences

A world may contain a game

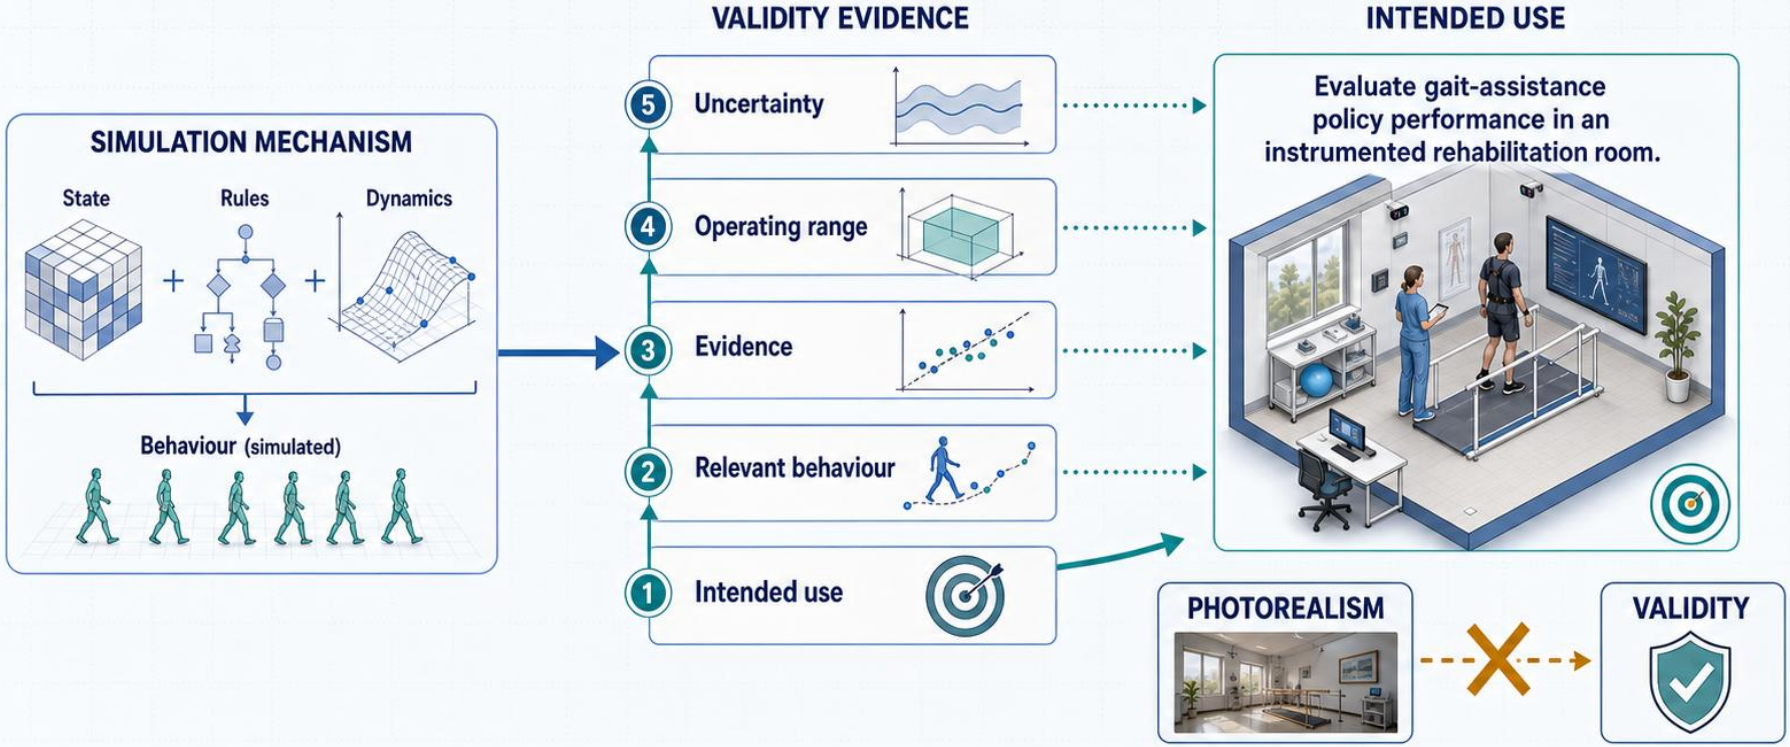
A game may use a thin environment

XR is an access and alignment system



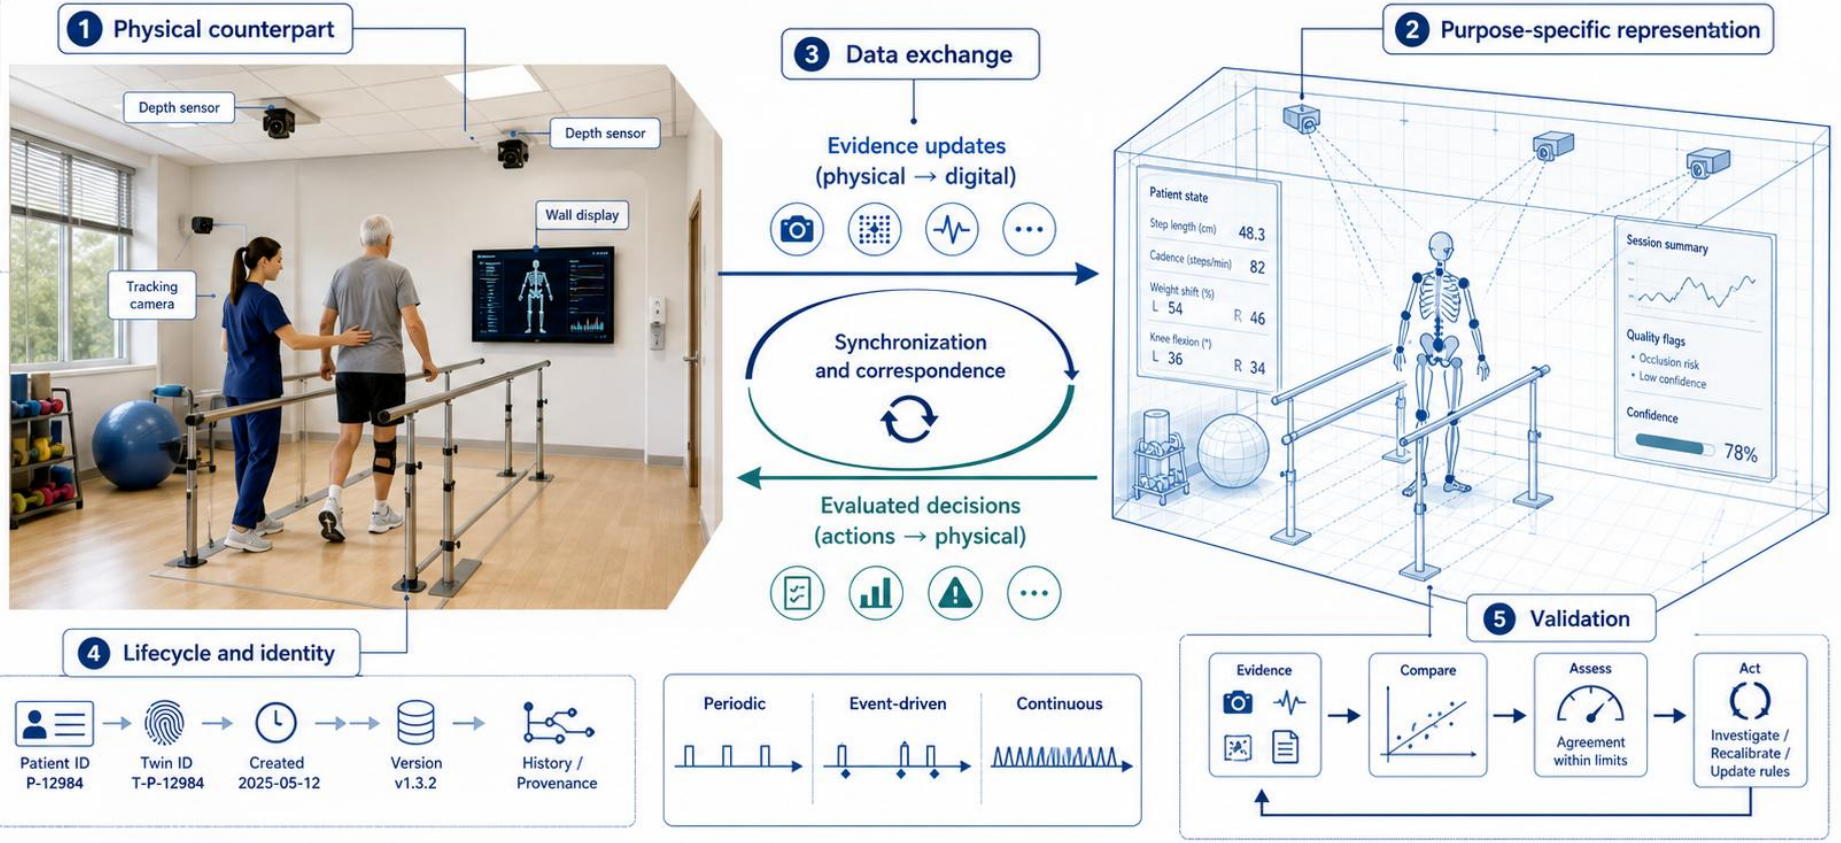
XR IS THE INTERFACE — THE VIRTUAL WORLD IS THE SYSTEM.

Simulation makes a validity claim



Appearance may help — validity requires task-relevant evidence.

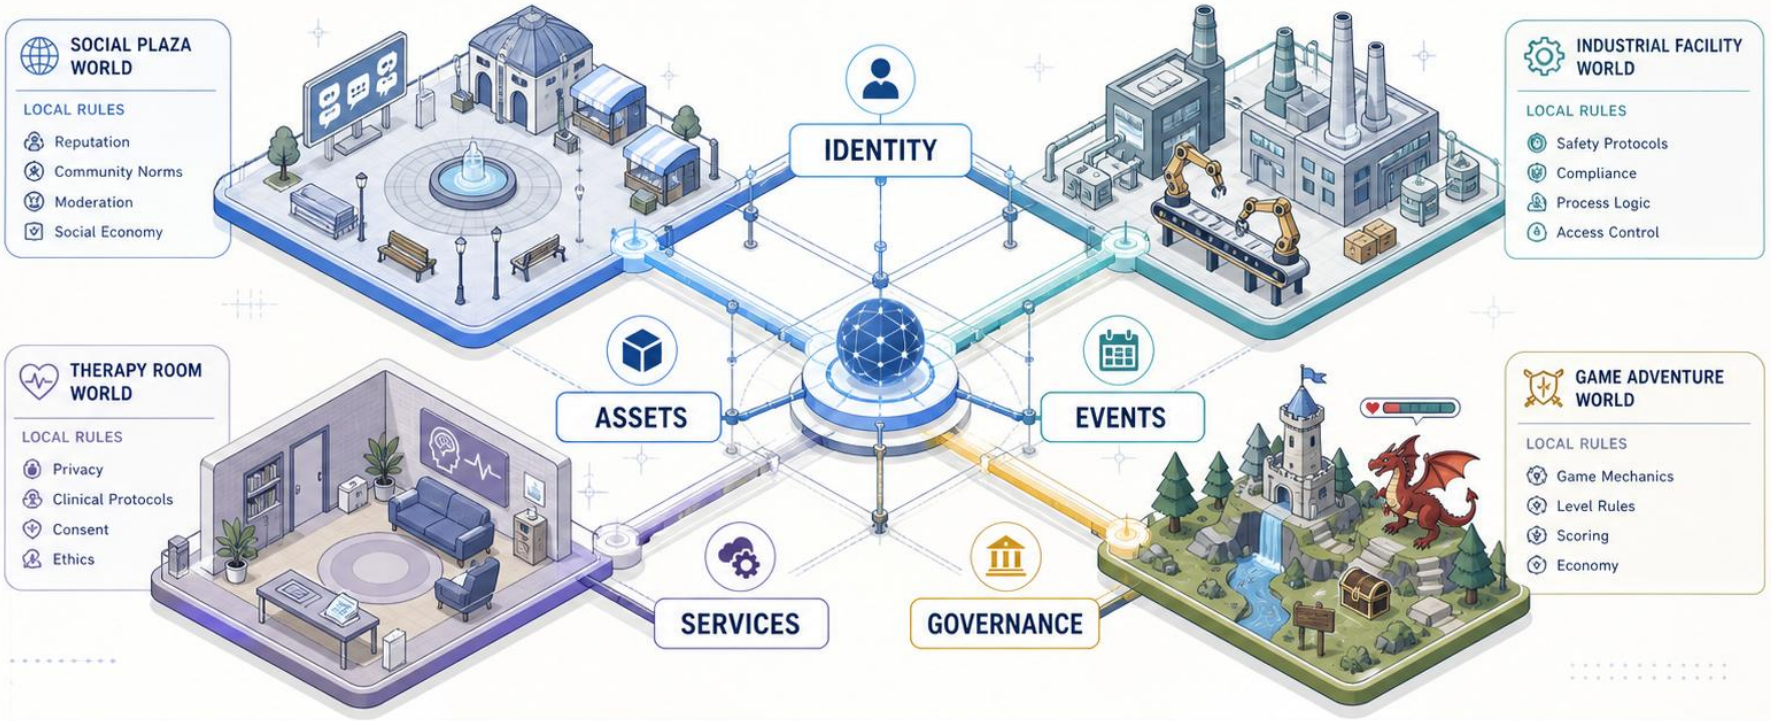
A digital twin is a maintained relationship



Lesson 03 — acquisition, synchronization, twin architecture

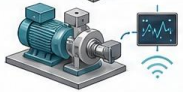
This is a preview only, not a claim that visual similarity guarantees validity.

Metaverse proposals add interoperability



SEMANTIC | TECHNICAL | LEGAL | GOVERNANCE

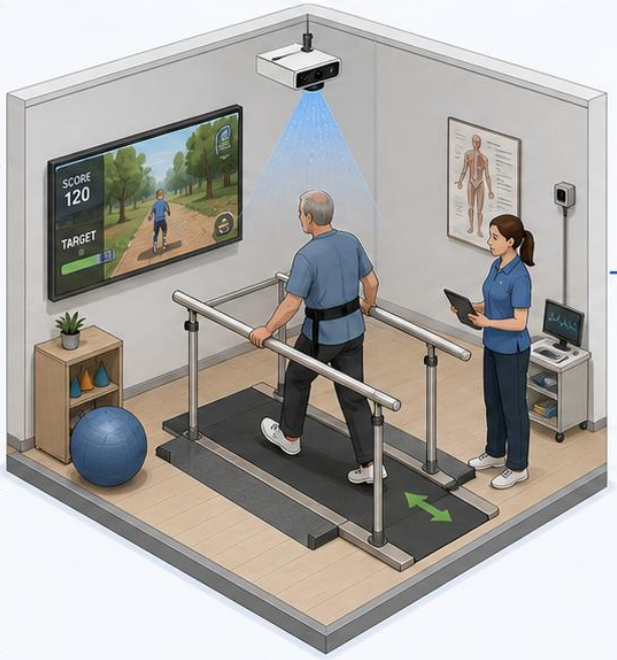
What can we use virtual worlds for?

USE	CONCRETE EXAMPLE	WHAT THE WORLD ENABLES
 **Play and socialize**	 Cooperative quest or live virtual event	 Shared state, identities and interaction
 **Learn and explore**	 Virtual science lab or historical city	 Change variables; inspect consequences
 **Rehearse procedures**	 Spacewalk or emergency-response drill	 Repeat rare situations; debrief actions
 **Design and plan**	 Factory layout or building waikthrough	 Compare alternatives before construction
 **Monitor operations**	 Sensor-linked equipment twin	 Relate live state to expected behavior
 **Train autonomy**	 Driving or indoor-navigation scenarios	 Practice policies under controlled variation
 **Generate labeled data**	 Vary lighting, objects and viewpoints	 Export images, depth, masks and poses
 **Test intelligent agents**	 Instruction-following tasks in a room	 Measure completion, errors and recovery



Choose state, rules and fidelity for the task. Validate transfer to reality.

Virtual worlds give AI controlled experience



CONTROLLABLE FACTORS

RESET

DUPLICATE

ACCELERATE

×2

INSTRUMENT

VARY



MEASUREMENTS

state

actions

events

OUTCOMES

train

test

causal probes

safety

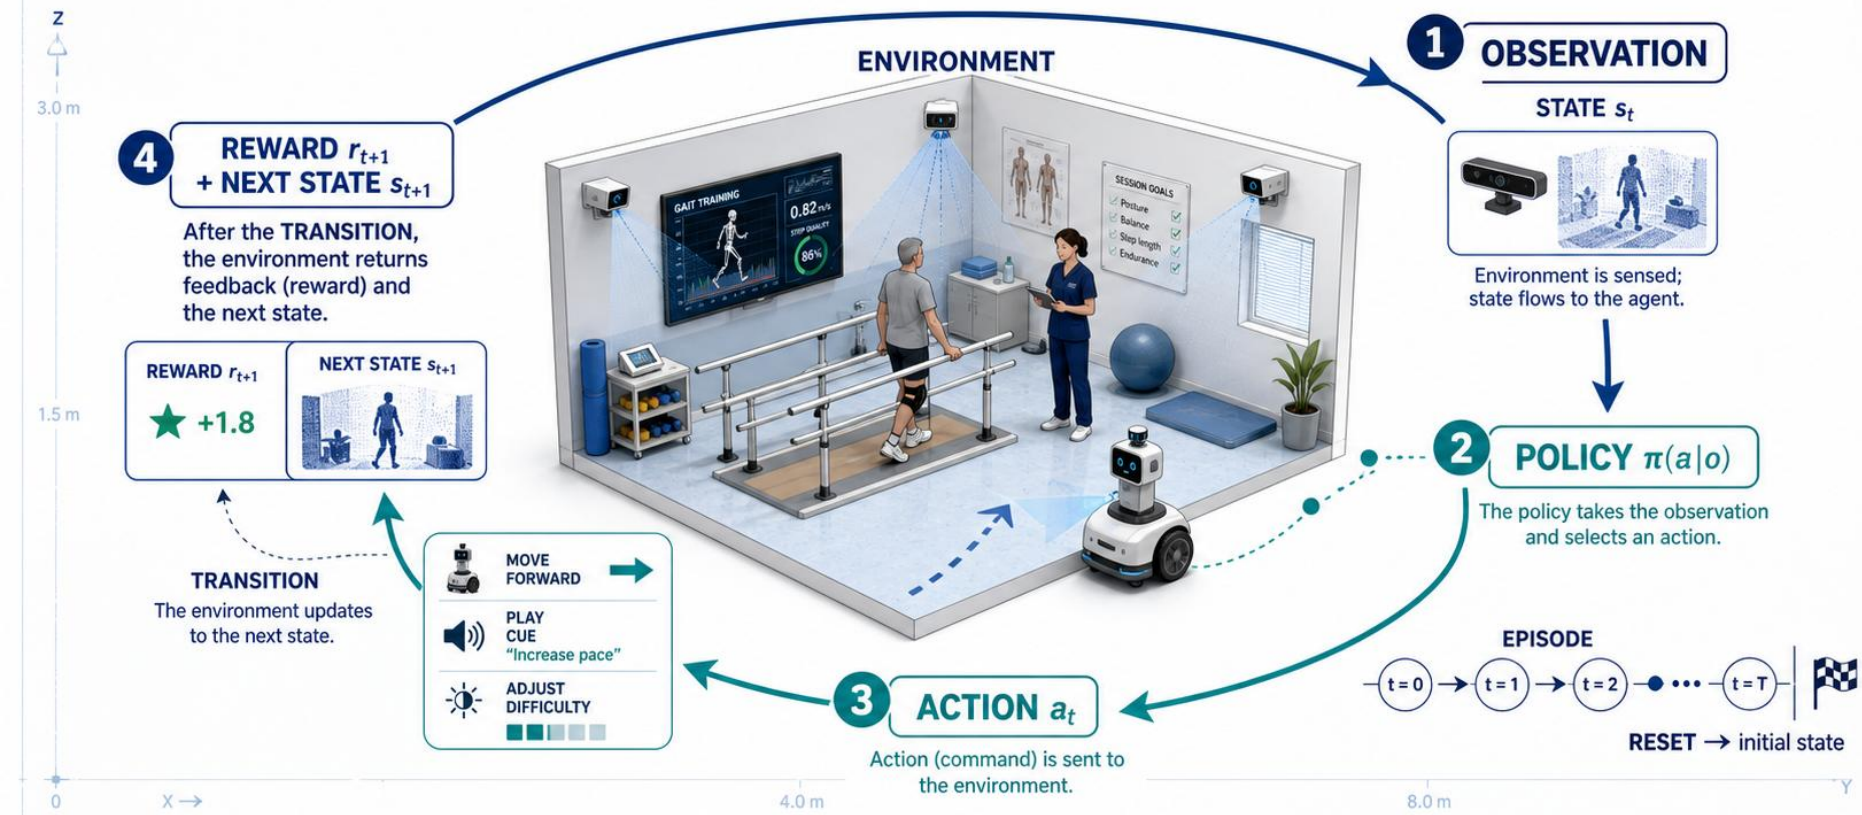
RECORDER

STATIC DATASET: fixed observations

VIRTUAL WORLD: actions → consequences

VALIDITY DEPENDS ON ASSUMPTIONS + COVERAGE

Reinforcement learning closes the loop

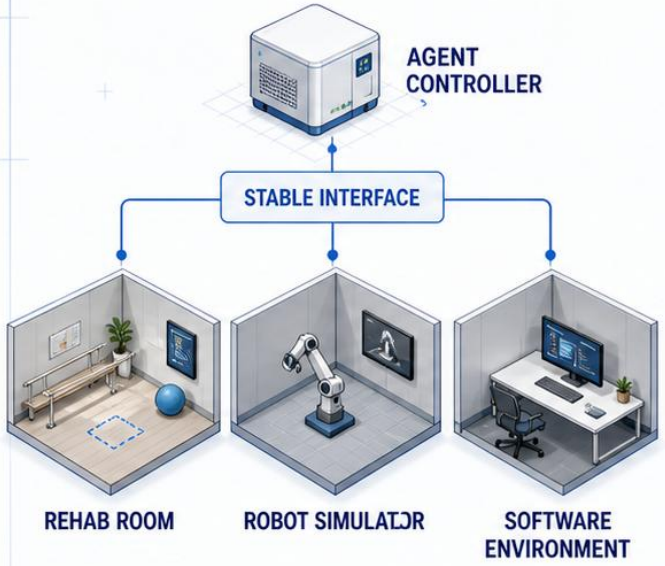


Reward is a training signal — not a complete specification of desired behaviour.

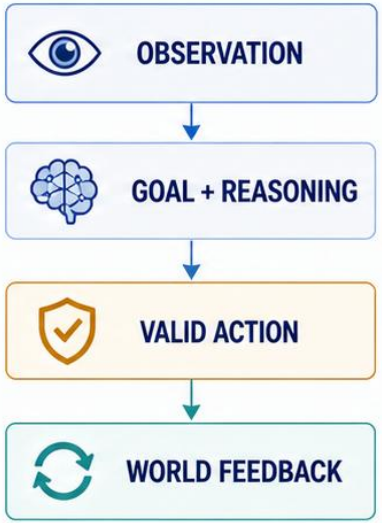
Lessons 06–07

Language and multimodal agents need grounded action

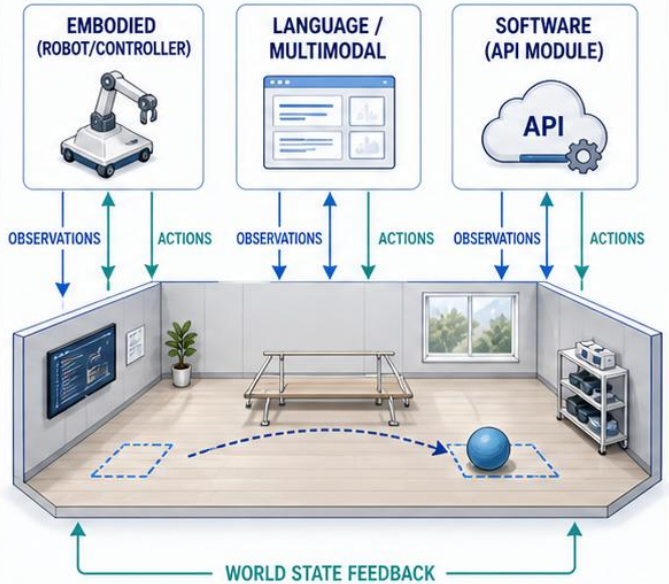
ONE AGENT / MANY WORLDS



GROUNDED-ACTION CONTRACT



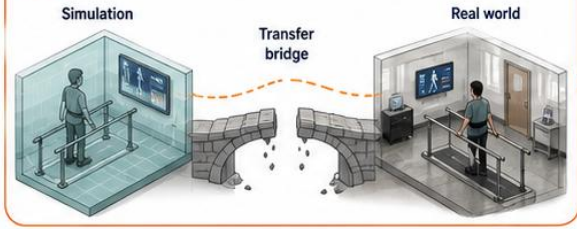
MANY AGENTS / ONE WORLD

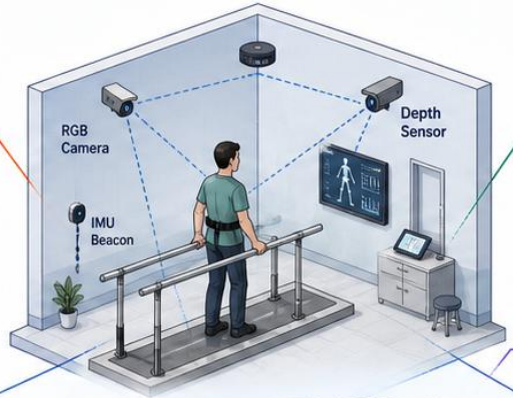




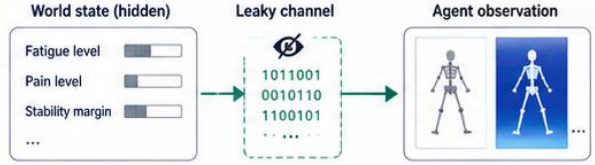
Worlds can teach the wrong lesson

1. Sim-to-real gap

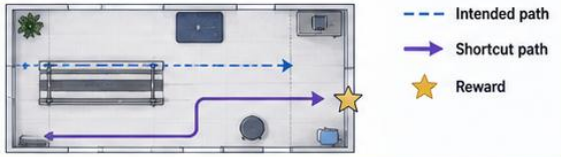




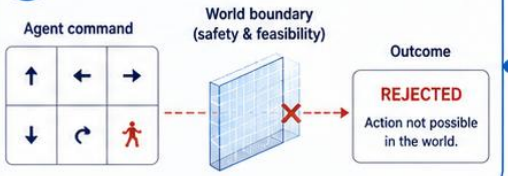
2. Hidden-state leakage



3. Reward shortcuts

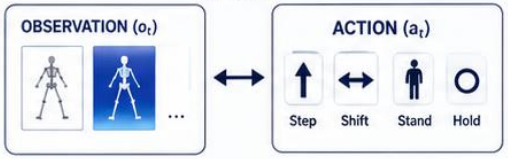


4. Invalid actions

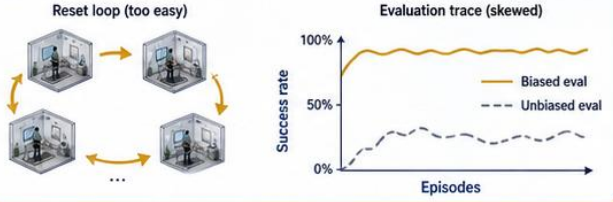


AGENT ↔ WORLD

Interface



5. Biased resets and metrics

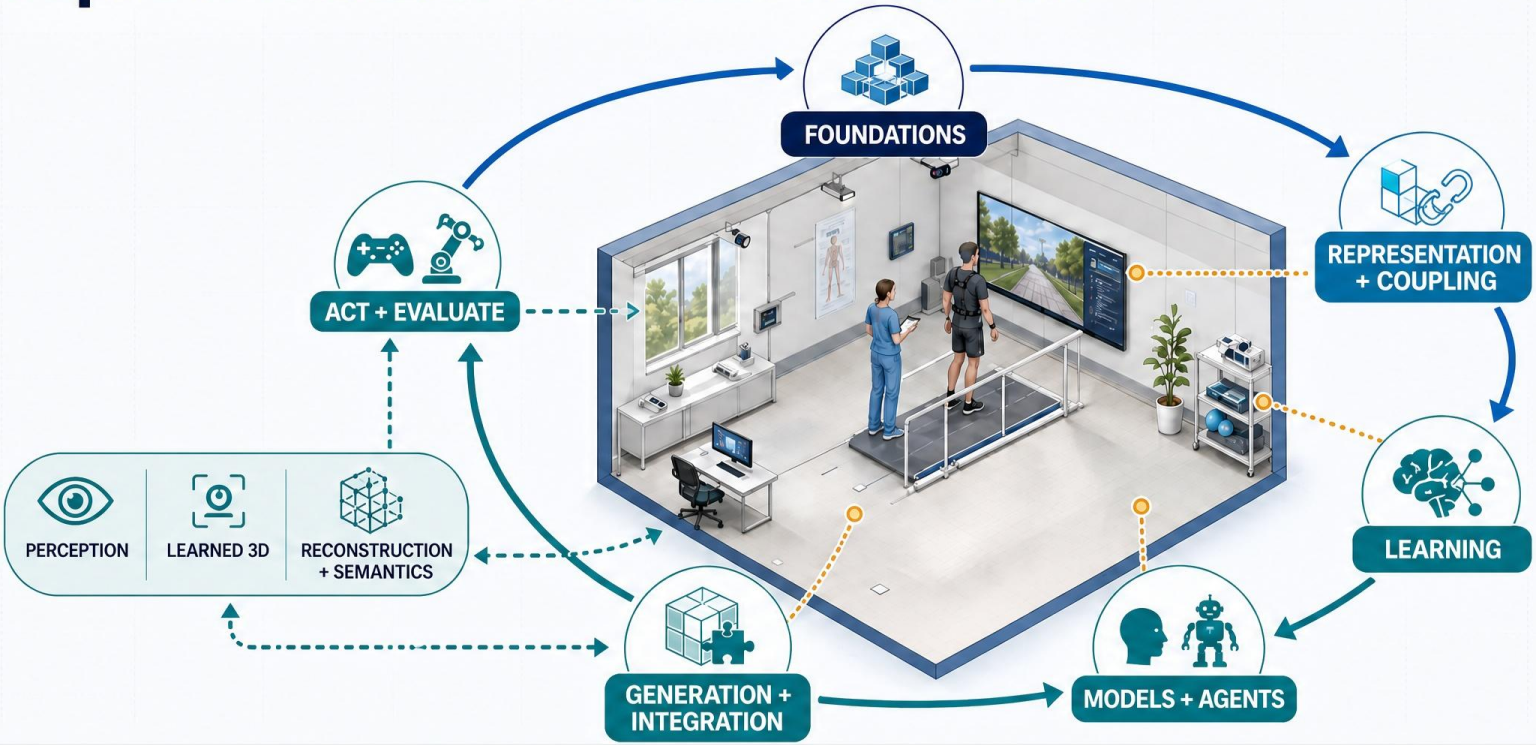


DIAGNOSTIC TRACE

1. Sim-to-real gap
2. Hidden-state leakage
3. Reward shortcuts
4. Invalid actions
5. Biased resets and metrics

WHAT DID THE WORLD MAKE EASY?

The course connects representation to situated action



Build the world → couple perception → learn and reason → act → evaluate

One instrumented room accumulates every layer



GEOMETRY
Mesh, surfaces and spatial structure

STATE
Pose, objects and live digital twin

OBSERVATION
Sensors, XR view and robot perception

ACTION
Commands from phone, robot, RL, clinician

PHYSICAL COUPLING
Bidirectional loop between physical activity and twin

update

sense

act

update

Physical world (activity & sensors)

Digital twin (state & model)

PHONE / XR
observe
act

ROBOT
observe
act

RL POLICY
STATE → POLICY $\pi(a|s)$ → ACTION

LANGUAGE AGENT

```
User: increase speed slightly
Agent: calling set_ speed(1.1)
>
```

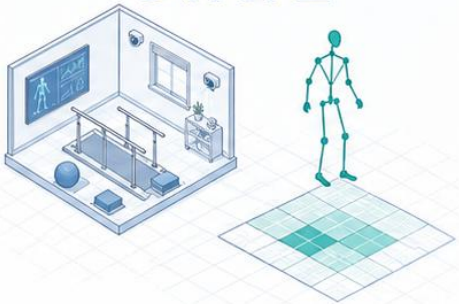
observe
act

SIMULATION + DATA

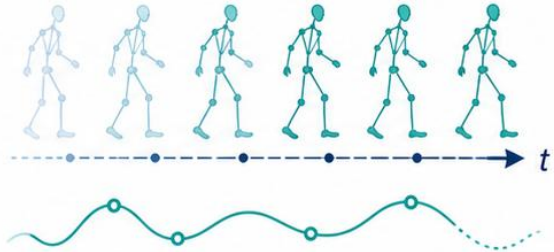
episode	trajectory	metrics
001	$[s_0, a_0, s_1, \ldots]$	⋯
002	$[s_0, a_0, s_1, \ldots]$	⋯
⋯	⋯	⋯

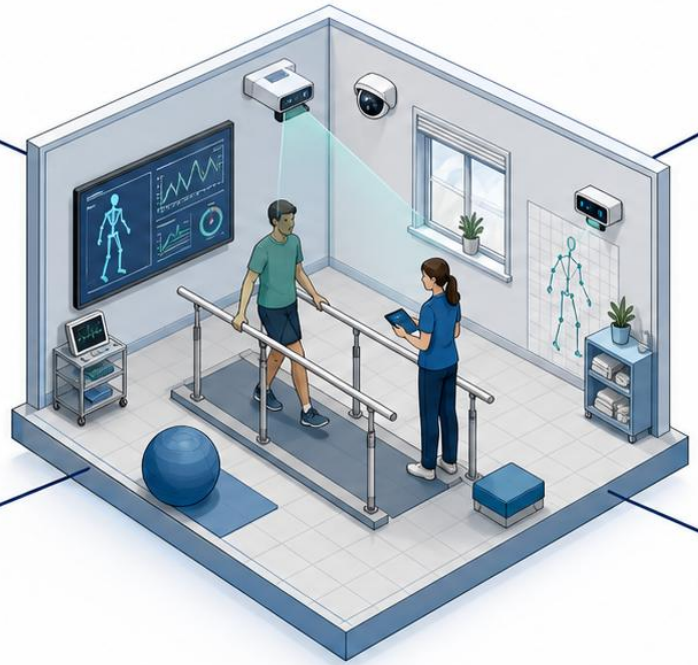
State, dynamics, agency, consequences

STATE

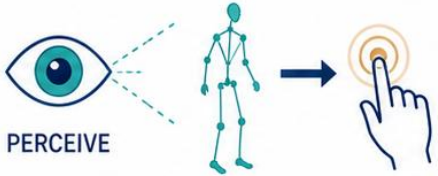


DYNAMICS





AGENCY



CONSEQUENCES



EXIT QUESTION: Which dimension would you inspect first when an apparently intelligent agent fails?

NEXT: COMPUTER GRAPHICS

References and further reading





DEFINITIONS

Bell, Toward a Definition of Virtual Worlds (2008)

Girvan, What is a virtual world? (2018)



STANDARDS & DIGITAL TWINS

W3C WebXR Device API

NIST digital twin overview

IEEE



AI & AGENTS

Sutton and Barto, Reinforcement Learning: An Introduction (2018)

Park et al., Generative Agents (2023)

SIMA Team (2024)



COURSE SOURCES

Complete citations: Lesson 01 overview and source ledger

▶ **Definitions:** Bell (2008); Girvan (2018) | ▶ **Standards:** WebXR; NIST digital twins; IEEE | ▶ **AI:** Sutton & Barto; Generative Agents; SIMA